Example Fair Machine Learning

Mortgage Lending Contest Rules

DISCLAIMER: These are sample proposed contest rules as might appear on a state informational website. They are intended to illustrate how a banking regulator in a state such as New York could pilot a fairness metric for machine learning based algorithmic lending by running a fair lending model contest in partnership with a university.

For more on the background for these sample contest rules, read the full automated fair lending proposal.

OVERVIEW

What is this contest?

Mortgage lenders increasingly use machine learning (ML) algorithms to make loan approval and pricing decisions. While these algorithms have the potential to expand access to credit by finding new creditworthy borrowers, careful oversight is needed to ensure they do not unfairly discriminate. Algorithmic lending brings a new set of challenges to fair lending assessments.

In order to pilot several innovative fair lending assessment methods, the state banking regulator (e.g., the New York Department of Financial Services) has announced an open Fair Machine Learning Lending Model contest. This contest is a partnership between the state banking regulator and state-based universities. The contest will pilot a new statistical fairness metric and a suite of automated fair lending tests that may be used by state banking regulators in the future to help detect and correct for bias in lending algorithms.

What do I have to do to enter?

We invite you to create a machine learning mortgage lending model optimized for a given definition of statistical fairness. Models should accept a borrower profile (including the loan amount and several measurements of creditworthiness) and output a loan decision (yes or no) and an interest rate.

We provide a training data set containing historic borrower profiles, loan decisions, and loan performance; a fairness definition; and a suite of automated tests that assess the model according to the fairness definition. You can test your model as many times as you like with the test suite. Once you're satisfied with your performance, submit your final model. At the end of the competition period, all models will be judged using a similar automated fairness test, but with different test data. This is to discourage models overfitting to the data in the provided automated test suite.

What's the prize?

The winning team will receive $5,000.

Fair Lending Contest Stages
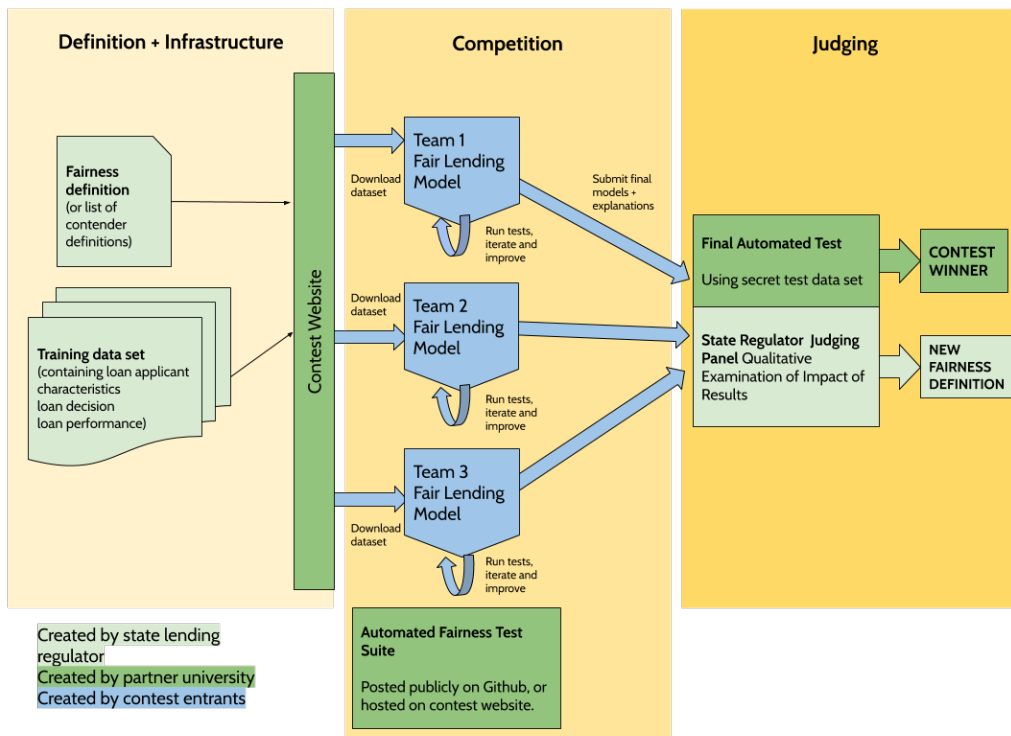
Who is eligible to enter?

Anyone who can build and train a machine learning model and is based in the United States is welcome to enter. Teams may include students, researchers, lending professionals, software engineers, designers, etc. Please keep in mind that you will need to submit your full code repository as a part of your entry. Thus, if you work in this field, you must enter in a personal capacity; do not include any proprietary company code or materials in your submission.

How long will the competition run?

You can download the training data set at the beginning of the competition stage. The suite of automated fairness tests will be made available at the same time. The competition will run for five months, during which time you can develop your model and run it through the fairness test suite as many times as you like. Once you are satisfied with your model's performance, upload it to the contest website for judging.

How will the competition be judged?

The winner is the team that creates an algorithm with the best performance on an automated fairness test similar to the one provided to you. The judging test will use a different set of data, so perfect performance on the test data does not necessarily guarantee perfect performance on the judging data. Judging will take place over the course of several weeks after the competition stage is over. The list of winners will be posted to the contest website.

The results from the contest will help state banking regulators decide whether the proposed fairness definition achieves their policy goals, and whether they should use this fairness definition to evaluate lenders' machine learning models for compliance with fair lending regulations.

DATASET USAGE LICENSE

The contest grants participants a limited, non-exclusive, royalty-free, non-transferable, and non-assignable license to reproduce, extract, or re-use the State Fair Lending Dataset (the "Dataset"), in whole or in part, for your research and evaluation purposes only. The license does not grant the right to distribute or sell the Dataset to any other party. For clarity, you may not distribute or provide this Dataset directly to another party for their use. Any third party must obtain the Dataset solely through this contest and agree to the License's terms.

(adapted from FICO dataset rules)

DATASET DETAILS

The contest materials include a baseline dataset for training and testing. The dataset includes the following data points for historical loans made in the United States:

‣ Amount of loan

‣ Loan-to-value ratio

‣ Debt-to-income ratio

‣ Loan approval or denial

‣ Loan interest rate

‣ Whether the loan was delinquent for more than a certain number of weeks over the course of its life.

Details on race, ethnicity, gender, or any other protected class variables will be excluded from the training data set. The test data used to score your model will include this information in order to calculate approval rates for each demographic group. This models the real-world scenario in which lenders collect protected-class information for HMDA reporting, but because of the disparate treatment standard are not allowed to use this data in making loan decisions.

CHALLENGE RULES AND SCORING

What are the submission requirements?

1. Submit a pull request to our GitHub repository with your model.

   a. The automated fairness tests must pass at least a certain bar in order for your submission to be considered.

   b. You can update your pull request and re-run the tests as many times as you like before the deadline. All submissions will be auto-scored at the time of the deadline.

2. Submit an explanation of your model, including answers to the following questions:

   a. What type of model did you use?

      i. Examples include linear regression, random forest, lasso regression, etc.

   b. Which demographic group does your model help the most (which group is most likely to get loans)?

   c. Which demographic group is least likely to get loans?

The winner will be the team with the model that performs best for the given fairness definition.

The test data used to decide the contest winners will be different from the quiz data used in continuous integration tests to give you progress updates. This is to prevent submitted models from overfitting to the quiz data set.

FAQ

When can I download the data?

The training data set and the fairness testing suite will be made available on the contest website and GitHub at the beginning of the contest period.

How long do I have to work on the model?

Five months from the start date of the competition phase.

When will I find out if I was chosen as the winner?

At the end of the competition period, there will be a judging period. After this, contest winners will be posted on the contest website and notified via email.

What data do I need to provide to enter?

The names and emails of everyone on your team so we can contact you, and the pull request to our Github repository containing your model. Your entire trained model must be in the pull request; your code may not make calls to any separate online services in order to make loan decisions.

Do I have to be a resident of a particular state to enter?

No, this contest is open to teams from anywhere in the United States.